

# Human-Centric Open-Future Task Discovery: Formulation, Benchmark, and Scalable Tree-Based Search

Zijian Song<sup>1</sup>, Xiaoxin Lin<sup>1</sup>, Tao Pu<sup>1</sup>, Zhenlong Yuan<sup>4</sup>, Guangrun Wang<sup>1,2,3\*</sup>, Liang Lin<sup>1,2,3</sup>

<sup>1</sup>Sun Yat-sen University

<sup>2</sup>Guangdong Key Laboratory of Big Data Analysis and Processing

<sup>3</sup>X-Era AI Lab

<sup>4</sup>AMAP, Alibaba

{songzj8,linxx76}@mail2.sysu.edu.cn, putao537@gmail.com

yuanzhenlong.yzl@alibaba-inc.com, wanggrun@gmail.com, linliang@ieee.org

## Abstract

Recent progress in robotics and embodied AI is largely driven by Large Multimodal Models (LMMs). However, a key challenge remains underexplored: how can we advance LMMs to discover tasks that assist humans in open-future scenarios, where human intentions are highly concurrent and dynamic. In this work, we formalize the problem of Human-centric Open-future Task Discovery (HOTD), focusing particularly on identifying tasks that reduce human effort across plausible futures. To facilitate this study, we propose HODT-Bench, which features over 2K real-world videos, a semi-automated annotation pipeline, and a simulation-based protocol tailored for open-set future evaluation. Additionally, we propose the Collaborative Multi-Agent Search Tree (CMASST) framework, which decomposes complex reasoning through a multi-agent system and structures the reasoning process through a scalable search tree module. In our experiments, CMASST achieves the best performance on the HODT-Bench, significantly surpassing existing LMMs. It also integrates well with existing LMMs, consistently improving performance.

## Introduction

Developments in robotics and embodied AI hold great promise for assisting humans in daily life. Recent advancements in Large Multimodal Models (LMMs) have significantly accelerated this process, empowering robots with remarkable intelligence in various domains (Khandelwal et al. 2022; Liang et al. 2023; Lin et al. 2023b; Driess et al. 2023; Huang et al. 2023; Yu et al. 2023). Most recently, research has begun to leverage Large Multimodal Models (LMMs) to enable robots to autonomously acquire new skills and experiences in unseen environments, a concept known as Autonomous Skill Acquisition (Zhou et al. 2024; Ahn et al. 2024; Yang et al. 2024b; Katara, Xian, and Fragkiadaki 2024; Bharadhwaj et al. 2024).

A central capability for autonomous agents is *task discovery*, where LMMs propose manipulation tasks for robots to execute (Wang et al. 2023b; Ahn et al. 2024; Yang et al.




Figure 1: **The illustration of HODT.** Driven by an overall goal, humans often engage in concurrent sub-processes, resulting in multiple possible future branches. HODT aims to identify tasks that remain helpful across diverse and uncertain futures. For example, as highlighted by the green box, completing *wipe the table* in advance saves human effort regardless of the order of other steps.

2024b; Bharadhwaj et al. 2024). While recent methods focus on generating tasks based on current observations, they typically assume fixed goals or closed environments. However, real-world human contexts are far more complex: people often engage in multiple sub-processes simultaneously, shift intentions dynamically, and rarely make all future steps explicit. This gives rise to the critical problem of *Human-Centric Open-Future Task Discovery*—inferring tasks that remain helpful across diverse, uncertain future trajectories. Unlike traditional task discovery, which aims to find the next best step toward a known outcome, open-future discovery must anticipate a range of plausible futures and identify actions that support them all. For example, as illustrated in Fig. 1, a robot assisting with housework should proactively wipe the table—a task that remains useful whether the human later cooks, cleans, or rests. Solving this problem is essential for enabling robots to provide anticipatory, generalizable support in dynamic, human-centered environments. It marks a necessary step toward collaborative AI that is not just responsive, but truly aligned with human intent.

Given these important but neglected aspects, we introduce the novel problem of **Human-centric Open-future Task Discovery**. To clearly study it, we begin by formally defining the HODT problem, motivated by its central premise: dis-

\*corresponding author

covering human-centric tasks that arise from open-ended future possibilities. A formal definition of human-centric tasks is also provided to support this formulation.

However, assessing task-discovery performance under the open-future setting poses significant challenges. As discussed above, the concurrent nature of human behavior leads to exponential growth in possible future branches, making it prohibitively expensive to annotate the complete set of helpful tasks. Moreover, relying on human annotators may introduce subjective bias. To address these issues, we present the HOTD-Bench along with a simulation-based evaluation approach. The HOTD-Bench is sourced from existing datasets and encompasses diverse real-world activities. The simulation evaluation adopts a discriminative strategy, avoiding exhaustive enumeration while ensuring alignment with human preferences. By leveraging the world-knowledge in LLMs (Jin et al. 2024; Gruver et al. 2023; Cao et al. 2024), our evaluation method accommodates open-set hypothetical future branches, including those not realized in the observed scenario. This mitigates the limitations of purely observational evaluation, enabling comprehensive assessment of both practical and theoretically optimal task decompositions. Experimental results on the HOTD-Bench reveal that existing LLMs achieve only limited performance on HOTD.

To bridge this gap, we introduce a Collaborative Multi-Agent Search Tree framework (CMAST). Our core idea is to construct the search tree of procedural structure and identify appropriate tasks accordingly. Our framework has two major innovations. First, CMAST introduces a search tree module to explicitly structure the reasoning process. By iteratively building the search tree, our model captures the inherent uncertainty of action procedures, thoroughly exploring various future scenarios. Furthermore, the search tree module allows for a scalable test-time thinking, a key trait shared with OpenAI-O3 (OpenAI 2025) and DeepSeek-R1 (Guo et al. 2025). Second, CMAST employs a collaborative multi-agent system, where specialized agents manage different stages of the reasoning. This collaboration effectively decomposes the complex reasoning, enabling each agent to focus on a specific aspect, reducing difficulty.

Experimental results show that our framework significantly outperforms existing LLMs in terms of Valid Task Ratio, while maintaining competitive performance in Valid Task Count, demonstrating its strong advantage in the HOTD. Ablation studies confirm the effectiveness of the search tree module and demonstrate that our framework can seamlessly integrate with various LLMs. Visualizations further illustrate the framework’s ability to suggest appropriate tasks by exploring diverse future procedures. Additionally, experiments validate the effectiveness of the simulator in reasonably deducing future scenarios.

Our key contributions are as follows: (1) We introduce and formulate the Human-centric Open-future Task Discovery problem, contributing to effective human-AI collaboration. (2) We establish HOTD-Bench, consisting of over 2K real-world videos from two sources. A simulator is proposed to deduce a given task’s future outcome, enabling the assessment of any future trajectory and the evaluation of its helpfulness. (3) We propose the CMAST framework to man-

age the complex reasoning. Experiments show that CMAST framework can seamlessly integrate a variety of existing LLMs and achieve consistently superior performance.

## Related Work

**Autonomous Skill Acquisition.** This technology encourages robots to learn new skills with less human instruction (Bharadhwaj et al. 2024; Ahn et al. 2024). Some studies focus on simulation-based learning (Wang et al. 2023b,a; Zhao, Weber, and Wermter 2024; Katara, Xian, and Fragkiadaki 2024; Yang et al. 2024b), leveraging LLMs to automatically generate simulation components, such as scene assets and supervision. Other studies focus on real-world learning (Ahn et al. 2024; Zhou et al. 2024). They typically leverage LLMs to recommend physical-interactive tasks, allowing the robot to attempt them and gather experience. In this work, we aim to further enhance LLM’s task-proposing capabilities, focusing on generating more valuable tasks to assist humans. Through this enhancement, we hope robots will acquire skills that directly align with human needs.

**Enhancing LLMs for Complex Reasoning.** Recent studies have focused on enhancing LLMs for complex reasoning, with two common approaches being Multi-Agent Systems and Chain-of-Thought. Multi-agent system deploys multiple agents to break down complex problems into smaller, manageable sub-parts (Wu et al. 2023; Hong et al. 2023; Yu et al. 2025), which has been verified across various applications (Yang et al. 2024d,c; Aher, Arriaga, and Kalai 2023). Chain-of-thought reasoning enhances LLMs by generating intermediate steps that progressively lead to the final answer (Guo et al. 2025). The most notable models are DeepSeek-R1 (Guo et al. 2025), OpenAI-o1/o3 (OpenAI 2024, 2025), which feature a *scalable test-time thinking* that performs increasingly efficacious CoT reasoning with extended thinking time. Inspired by previous works, our approach combines the advantages of both approaches.

**Related Video Datasets.** Video understanding has always been a research focus. There exist a number of video understanding datasets (Soomro 2012; Caba Heilbron et al. 2015; Jia et al. 2020; Damen et al. 2022; Zeng et al. 2024; Zhou et al. 2023), some of which are related to humans (Zellers et al. 2019; Lei et al. 2018; Grauman et al. 2022). However, none of the existing datasets focus on task discovery, where the answers are supportive tasks for humans. In this work, we curate our benchmark based on the existing video datasets and develop an evaluation approach to assess the contribution of any discovered task.

## Formulation

### Problem Formulation

The HOTD problem aims to discover a set of tasks that assist the human in a given video segment. Specifically, given an input video segment  $I_{0:t_0}$ , the model generates a set of predicted tasks, denoted as  $\hat{Q}_I = \{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_i\}$ , where  $\hat{\cdot}$  indicates the predicted value.

To formally define the objective of HOTD, we denote  $Q$  be the set of all task descriptions in natural language. For a given video input  $I$ , we define  $Q_I^{hc}$  as the subset of  $Q$  that

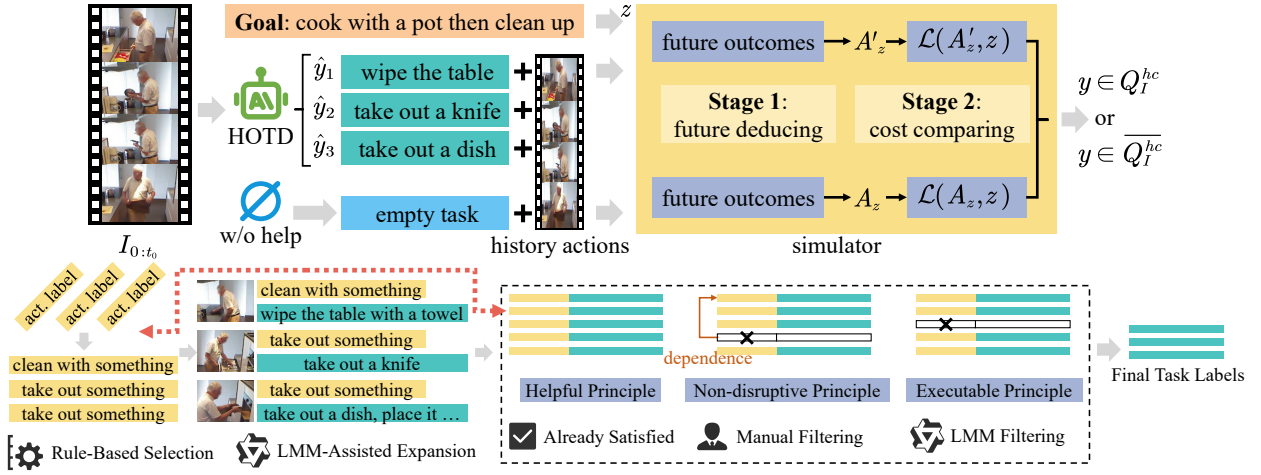


Figure 2: **The process of simulation-based evaluation (upper) & annotation pipeline (lower).** The simulator first takes a discovered task  $\hat{y}_n$ , a historical action sequence and its associated goal  $z$  as input. Then, it simulates the resulting future trajectory  $A'_z$  by accounting for the adjusted human actions until the goal. Finally, it summarizes the overall process to estimate the incurred cost  $\mathcal{L}(A'_z, z)$  and compares it to the original cost  $\mathcal{L}(A_z, z)$ . In the annotation pipeline, future actions are first selected to meet the *helpful principle*, then expanded into descriptive sentences and filtered through the *non-disruptive principle* and the *executable principle*, finally forming the task labels.

contains only the human-centric tasks constrained by the observed scenario  $I$ . An HOTA method is characterized by a mapping function  $G$ , which predicts a set of tasks based on the given video input:  $\hat{Q}_I = G(I)$ . The objective of HOTA is to find an optimal mapping function  $G$  that maximizes the inclusion of human-centric tasks:

$$\max_G |\hat{Q}_I \cap Q_I^{hc}|, \quad \max_G \frac{|\hat{Q}_I \cap Q_I^{hc}|}{|\hat{Q}_I|}, \quad (1)$$

where the second is the normalized version of the first.  $\hat{Q}_I$  denotes the predicted task set,  $Q_I^{hc}$  denotes the ground truth human-centric task set. We further use  $\overline{Q_I^{hc}}$  to denote the complement of  $Q_I^{hc}$ , only containing unhelpful tasks.

For the notion of open-ended future possibilities, we do not explicitly define it. Instead, we rely on the simulation evaluation, which accommodates open-set future scenarios and estimates outcomes in line with human preferences.

## The Definition of Human-Centric Tasks

To establish a clear understanding of HOTA, we develop a systematic definition of what human-centric tasks are. Concretely, a human-centric task is defined as an executable action that contributes to achieving the human’s goal.

First, each human-centric task corresponds to a specific action that can be executed by robots. We follow the three-level taxonomy, ‘action primitive’, ‘action’, and ‘activity’, introduced by Moeslund et al. (Moeslund, Hilton, and Krüger 2006). The human-centric task are defined to be at the second level, ensuring that they are neither too narrow to lack meaning nor too broad to confuse the robot.

Second, a human-centric task is not merely an action but one that assists a human in reaching their goal. Let

$z \in Z$  denote a human’s latent goal. The human executes a sequence of actions to accomplish  $z$ , denoted as  $A_z = \{a_1, a_2, \dots, a_n\}$ . When the robot performs an additional task  $y$ , the human may adjust their actions accordingly, resulting in a modified action sequence  $A'_z = \{y, a'_1, a'_2, \dots, a'_n\}$ , where  $a'_n$  represents the human’s modified action due to the influence of the robot’s intervention.

To formally define whether  $y$  provides assistance, we introduce a cost function  $\mathcal{L}$  that quantifies the cost required to achieve  $z$ . This cost function may be defined in terms of time spent, labor exerted, or other measures. A task  $y$  is considered helpful if its inclusion reduces the overall cost, where the subscript  $I$  denotes the history video driven by  $z$ :

$$y \in Q_I^{hc} \iff y \in Q \wedge \mathcal{L}(A'_z, z) < \mathcal{L}(A_z, z). \quad (2)$$

Conversely, a task that increases the cost is not helpful:

$$y \in \overline{Q_I^{hc}} \iff y \in Q \wedge \mathcal{L}(A'_z, z) \geq \mathcal{L}(A_z, z). \quad (3)$$

Such a discriminative definition is particularly appropriate and necessary, as the complexity of human cognition makes a prescriptive one infeasible.

## Benchmark

### Data Collection

In order to evaluate the model’s performance under practical scenarios, we construct our dataset from two existing datasets: the Toyota Smarthome Untrimmed (TSU) (Das et al. 2019; Dai et al. 2022) and the Charades (CHA) (Sigurdsson et al. 2016). They offer various real-world activities from two distributions. We apply a sliding window to segment the videos and filter out low-quality samples, resulting in 2450 curated clips totaling nearly 40 hours. The TSU and CHA subsets contribute 2K and 0.4K videos, respectively.

## Evaluate by Simulation

Under the open future setting, our evaluation aims to quantify how many human-centric tasks are included in the predicted set. However, such evaluation is nontrivial. A straightforward approach would be to have human annotators label all helpful tasks. While intuitive, it is impractical due to the annotator subjectivity and the prohibitive cost of exhaustively labeling exponentially many future branches (see Fig. 1). Instead, verifying whether a given task is helpful is substantially more tractable, as it only requires estimating the cost introduced in Eq. (2) and Eq. (3). This motivates our use of simulation as an evaluation tool. A simulator can flexibly model how the future would unfold under any hypothetical task insertion, thereby enabling the evaluation of arbitrary candidate trajectories.

The proposed simulation-based evaluation approach is illustrated in the upper part of Fig. 2. For sequences without robot intervention, it directly simulates the future processes. For sequences with robot intervention, it models human adaptation and reconstructs the complete sequence. The resulting trajectory is then used to estimate the incurred costs. A central strength of this approach lies in its generative nature, which allows for evaluating not only observed trajectories but also any hypothetical future beyond the dataset, thereby capturing scenarios that could be more optimal than those explicitly performed.

In our implementation, we adopt an LLM as the simulator, which has been proven to reliably deduce future evolution (Jin et al. 2024; Gruver et al. 2023; Cao et al. 2024), while also aligning with human preferences and minimizing subjective bias (Rafailov et al. 2023; Bai et al. 2022). The latent goal  $z$  is pre-annotated according to the whole video. The cost is defined in terms of time consumption. To mitigate the sensitivity caused by absolute time estimations, our simulator evaluates the relative time costs by comparing two action sequences with or without robot intervention. We adopt a chain-of-thought prompting approach that first guides the LLM to envision future outcomes and then determines whether  $\mathcal{L}(A'_z, z) < \mathcal{L}(A_z, z)$ . The experiment in Sec. 6.3 as well as Fig. 4 (a) and Fig. 5 validates that our LLM simulator is simple yet reliable.

## Evaluate by Labels

We introduce an alternative evaluation approach based on labeled task sets  $\tilde{Q}_I^{hc} \approx Q_I^{hc}$ , enabling a more stable evaluation. The core idea is to approximate a “ground truth” set of human-centric tasks, assuming that tasks humans will inevitably perform are inherently beneficial if completed in advance, as they directly reduce human effort and cost.

To construct  $\tilde{Q}_I^{hc}$ , we design a semi-automated annotation pipeline (see Fig. 2 lower). We begin with the dense action labels from the original datasets, select the future-performed actions, and expand them into descriptive, open-set task sentences using Qwen-VL (Wang et al. 2024). The annotation is guided by three predefined principles defining human-centric tasks. *Helpful*: A human-centric task should complete actions in advance that humans would otherwise have to do themselves. *Non-disruptive*: A human-centric

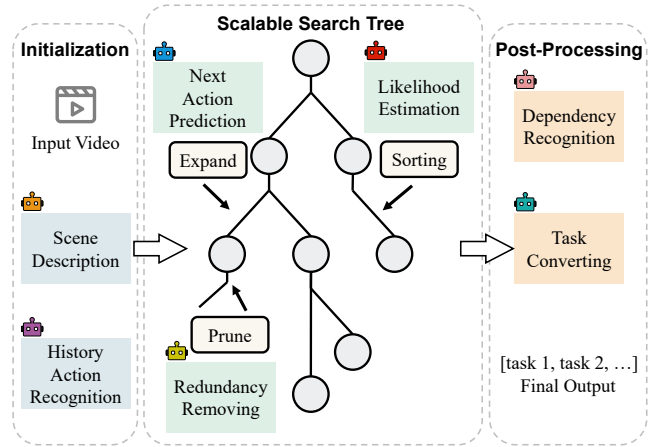


Figure 3: **The overview of the Collaborative Multi-Agent Search Tree framework.** It structures the HOTD reasoning with 7 LLM agents and a scalable search tree module.

task should not conflict with the human’s plan. *Executable*: A human-centric task must be possible to execute given the current conditions. The detailed implementation of them can be found in Appendix. By filtering and refining the expanded action descriptions under these principles, we obtain a reliable approximation of the ground truth task set.

## Method

### Search Tree Module

The HOTD reasoning not only requires understanding the visual contents, but also analyzing the open-future procedure. To achieve this, we build a search tree that explicitly represents open-future action space, enabling an explicit exploration of future scenarios. Additionally, it supports flexible expansion and pruning, facilitating *scalable test-time thinking* that achieves comprehensive answers, a key trait shared with OpenAI-O3 (OpenAI 2025) and DeepSeek-R1 (Guo et al. 2025).

The search tree is comprised of a set of nodes and edges:  $T = (V, E)$ . Each node corresponds to an individual action. The edge between nodes represents the temporal order.

The first  $N$  nodes in the tree represent the history action sequence determined by the input video, formulated as:

$$T^{0:N} = P(v^0, v^N) = (v^0, v^1, \dots, v^N); \quad (4)$$

$$\forall n < N, |Child(v^n)| = 1.$$

The superscript  $n$  represents the  $n$ -th layer of the tree. The symbol  $P(v, vt)$  represents the path from node  $v$  to node  $vt$ . The  $Child(v)$  represents the child nodes of the node  $v$ .

Starting from the  $N$ -th node, the tree begins to branch out. Each branch represents a possible next step, reflecting the uncertainty of the open-future, formulated as:

$$Child(v^n) = g(v^n) = \{v_1, v_2, \dots, v_k\}, n > N. \quad (5)$$

where  $g$  is a node expansion function which we will introduce in Eq. (6). As the tree progresses, multiple branches capture various possible action sequences. The leaf nodes in

the tree represent the completion of the entire activity, which we refer to as the ‘finish activity’.

The search tree supports several searching strategies. To balance performance and efficiency, we employ a pruned exhaustive search with a 0.5 probability threshold. Alternative strategies are discussed in the Sec. 6.3.

### Collaborative Multi-Agent System

Inspired by previous works (Wu et al. 2023; Hong et al. 2023; Yu et al. 2025), we develop a multi-agent system to structure the reasoning process in the HOTD problem. The workflow is presented in Fig. 3. Our key idea is to decompose the complex problem by aligning each agent with a specific stage of a search tree, such as initialization, expansion, pruning, and post-processing. This design not only preserves the generalization ability of LMMs but also enables seamless integration with various LMMs.

First, receiving the input video, the *Scene Description Agent* digests the video and produces a detailed description  $s = f_s(I_{0:t_0})$ , providing an overall precondition.

Second, the *History Action Recognition Agent* recognizes the history actions of the person,  $(v^0, v^1, \dots, v^N) = f_r(I_{0:t_0}, s)$ , which initializes the search tree.

Third, from the initial search tree, three agents are employed to expand the search tree iteratively. In each iteration, the *Next Action Prediction Agent* forecasts the next immediate action given the entire action path up to now. And the *Likelihood Estimation Agent* predicts the probability for each child node, providing reference for sorting and pruning. We manually add the ‘finish activity’ node to every non-leaf node. When the above two agents expand the search tree, the *Redundancy Removing Agent* is employed to prune the redundant nodes. These three agents work iteratively until all unexpanded nodes are leaf nodes or reach a maximum tree height. Let  $g$  denote the above three agents together, and the expansion process can be written as:

$$\{v_1, v_2, \dots, v_k\} = g(v^n, I_{0:t_0}, s), n > N. \quad (6)$$

Fourth, the search tree is formatted into a set of action sequences by traversing all the paths. The *Dependency Recognition Agent* is adopted to identify and exclude actions with prerequisite, retaining only those executable ones.

Finally we dismantle the sequential structure, yielding a set of independent actions. The *Task Converting Agent* is employed to transform each action into a task description from the robot’s perspective, formulated as:

$$\begin{aligned} \hat{Q}_I &= \{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_i\} \\ &= f_c(\{v|v \in \cup_{n>N} V^n\}, I_{0:t_0}, s), \end{aligned} \quad (7)$$

### Implementation Details

Our framework is entirely training-free, requiring no fine-tuning of the overall system or any submodule. This design allows seamless integration of various LMMs. Specifically, the *Scene Description Agent*, the *History Action Recognition Agent*, and the *Next Action Prediction Agent* are LMM agents, implemented with LLaVA-Next-Video (Zhang et al. 2024), the other three agents are LLM agents, implemented with Qwen-LM (Yang et al. 2024a).

## Experiment

### Experimental Setups

**Baselines.** We conduct a quantitative study on the HOTD-Bench to evaluate the HOTD capabilities of the following open-source LMMs: Qwen2-VL (Wang et al. 2024), Qwen2.5-VL (Team 2025), InternVL2 (Team 2024), InternVL2.5 (Chen et al. 2024), Video-LLaVA (Lin et al. 2023a), LLaVA-Next-Video (Zhang et al. 2024). These baselines are prompted with the video to recommend some assisting tasks.

**Evaluation Metrics.** We establish two evaluation metrics, Valid Task Count and Valid Task Ratio, corresponding to two objectives in Eq. (1). The Valid Task Count, denoted as ‘vc’, measures the average number of helpful tasks discovered, indicating the model’s ability to identify a diverse range of helpful tasks. The Valid Task Ratio, denoted as ‘vr’, measures the average proportion of helpful tasks within each prediction, indicating the reliability and the precision of the model’s outputs. Their formulations are shown below, where  $N$  denotes the number of samples:

$$vc = \frac{1}{N} \sum_{n=1}^N |Q_{I_n}^{hc} \cap \hat{Q}_{I_n}|, \quad vr = \frac{1}{N} \sum_{n=1}^N \frac{|Q_{I_n}^{hc} \cap \hat{Q}_{I_n}|}{|\hat{Q}_{I_n}|}. \quad (8)$$

For the majority of experiments in this paper, we report the results of both evaluation approaches (by simulation and labels). We also measured videos with different observation lengths, using ‘@’ for distinguishing. For example, ‘vc@40’ represents the metric given 40 sec of observations.

### Overall Evaluation

Tab. 1 shows the performance of our framework against baseline LMMs. Our method demonstrates significant advantages over baselines, validating its effectiveness.

For the Valid Task Ratio, our method clearly outperforms the others across observation lengths. Specifically, all existing LMMs achieve a relatively low Valid Task Ratio, from 29.1% to 62.4%, struggling to provide sufficient assistance, likely due to their training on dialogue instruction corpora, which is inadequate to capture the behavioral expectations of humans. In contrast, our approach surpasses the second-best method by 15% to 22% in the TSU subset and achieves the same level as the second-best method in the CHA subset. This validates that by forecasting future situations and imagining tasks in the future context, our method can discover more supportive tasks within each prediction round.

For the Valid Task Count, our method also achieves competitive results. Our method outperforms the second-best by an average of 7.6% in the TSU subset. This confirms that our method is capable of discovering a diversity of tasks, with a significant proportion of them being beneficial, largely due to the search tree module’s exploration of potential future scenarios and activity procedures.

### Further Analysis

**Comparing Among Existing LMMs.** Based on the results in Tab. 1, we further analyze the performance of ex-

Method	TSU subset				CHA subset	
	vc@20 / vr@20	vc@40 / vr@40	vc@60 / vr@60	vc@80 / vr@80	vc@20 / vr@20	vc@30 / vr@30
Qwen2VL-7B	2.86 / 44.3%	2.71 / 44.2%	<b>4.03</b> / 40.6%	2.77 / 42.0%	2.06 / 43.1%	2.08 / 48.1%
Qwen2.5VL-72B	2.45 / 48.8%	2.47 / 47.6%	2.62 / 42.6%	2.22 / 43.0%	3.01 / 40.9%	2.87 / 43.0%
InternVL2-8B	2.23 / 58.4%	2.51 / 61.0%	2.44 / 62.4%	2.79 / 59.8%	2.47 / 54.5%	2.28 / <b>58.1%</b>
InternVL2.5-26B	1.57 / 40.5%	1.48 / 42.4%	1.42 / 40.9%	1.38 / 40.5%	1.88 / 42.2%	1.74 / 43.3%
Video-LLaVA-7B	0.47 / 40.8%	0.45 / 41.5%	0.40 / 40.6%	0.46 / 42.1%	0.36 / 36.0%	0.29 / 29.1%
LLaVA-NV-7B	3.62 / 48.9%	3.34 / 50.2%	2.92 / 48.8%	3.41 / 51.5%	<b>6.20</b> / 54.1%	<b>6.02</b> / 50.0%
LLaVA-NV-34B	3.28 / 41.5%	<u>3.39</u> / 44.2%	2.69 / 45.0%	<u>3.41</u> / 42.6%	<u>3.55</u> / 40.8%	<u>3.18</u> / 42.7%
CMAST(ours)	<b>3.90</b> / <b>71.7%</b>	<b>3.83</b> / <b>71.9%</b>	<u>3.89</u> / <b>72.0%</b>	<b>3.86</b> / <b>71.8%</b>	2.73 / <b>55.5%</b>	2.82 / <u>53.6%</u>

Table 1: **The quantitative comparison in the HOTD-Bench, evaluated by simulation.** The ‘vc@’ and ‘vr@’ denotes the Valid Task Count and the Valid Task Ratio at a specific observation length, higher is better. The best are highlighted by **bold** and the second best are highlighted by underline. CMAST demonstrates significantly better performance against other methods.




Figure 4: (a) **Human evaluation of the simulator.** Columns indicate how many annotators rated each task as helpful, where 0 means all rated it unhelpful and 5 means all rated it helpful. The distribution shows strong agreement between the simulator and human preferences. (b) **Ablations on the search-tree module and component agents**, with 40-sec TSU observations, evaluated by simulator (left) and labels (right). (c) **Ablation on the search strategy** in the TSU subset, 40 sec observation length, evaluated by simulation. The ‘Efficiency’ is measured by the average number of expansions.

Method	TSU subset		CHA subset
	vc@40 / vr@40	vc@60 / vr@60	vc@20 / vr@20
Qwen2VL-7B	0.82 / 14.1%	0.82 / 13.9%	0.58 / <u>13.2%</u>
InternVL2-8B	0.50 / 16.0%	0.52 / 17.1%	<u>0.89</u> / 5.3%
Video-LLaVA	0.22 / <u>20.6%</u>	0.22 / <u>21.5%</u>	0.09 / 9.2%
LLaVA-NV-7B	<u>1.51</u> / 15.0%	<u>1.61</u> / 15.0%	<b>1.00</b> / 8.1%
LLaVA-NV-34B	1.22 / 13.1%	1.30 / 13.8%	0.87 / 10.9%
CMAST(ours)	<b>1.92</b> / <b>38.7%</b>	<b>1.98</b> / <b>39.3%</b>	0.79 / <b>15.6%</b>

Table 2: **The quantitative comparison in the HOTD-Bench**, evaluated by labels. CMAST also demonstrates significantly better performance against other methods.

isting LMMs. Notably, there is generally a trade-off between Valid Task Count and Valid Task Ratio. For instance, InternVL2-8B (Team 2024) achieves the highest Valid Task Ratio among baselines, but a low Valid Task Count, indicating overly conservative predictions. LLaVA-Next-Video-7B (Zhang et al. 2024) achieves the highest Valid Task Count among baselines, but a low Valid Task Ratio, indicating aggressive predictions with many invalid tasks. Moreover, larger models offer no clear advantage over smaller models. This indicates that scaling up does not necessarily enhance task discovery capabilities. Overall, current LMMs still face notable challenges in addressing the HOTD problem.

**Reliability of the Simulator.** To investigate *whether the simulator aligns with human preferences*, we conducted a human evaluation. We randomly selected 25 helpful and 25

unhelpful tasks as marked by the simulator, and presented them in mixed order to 5 annotators. Annotators independently judged whether each task is helpful. As shown in Fig. 4 (a), tasks judged helpful by the simulator were largely considered helpful by humans, and similarly for unhelpful tasks. The result shows that the simulator reliably reflects human preferences, offering accurate assessments.

**Case Study of the Simulator.** Fig. 5 presents future deductions generated by our simulator, demonstrating its strong generalization even for scenarios not explicitly observed in the dataset. For instance, in the last row, it successfully anticipates that the person will need to retrieve the cup as a consequence of the robot’s action, an event that wouldn’t have happened otherwise. These examples highlight the effectiveness of our LLM-based simulator, as it offers a reliable and comprehensive assessment of all hypothetical tasks.

**Ablation on the Search Tree Module.** We conduct an ablation study to examine the overall contribution of the search tree module. The results are shown in Fig. 4 (b). The ‘CMAS-main’ is our main version. The ‘CMAS w/o tree’ removes the entire search tree module, replacing it with an LMM agent that directly predicts future actions. As can be seen from the results, replacing the entire search tree module reduces the Valid Task Ratio by 37%. This suggests that a single LMM agent may produce confined prediction paths. On the contrary, the search tree module addresses it by providing an explicit and structured procedure space, allowing for a thorough examination of different action procedures.

Goal	Task	Deduction	Helpfulness
Prepare a beverage and clean the kitchen	Wipe down the countertop	The person continue to prepare a beverage .... Wiping down the countertop helps in organizing the space, making it easier ... to access items ....	True
Work on a laptop and prepare food	Preheat the oven	If the person plans to use the oven for cooking or baking, preheating it ... could help them achieve their goal more efficiently.	True
Read and work on a laptop	Pick up a remote to change the channel	The person's purpose is to read and work .... The task of pick up a remote ... <u>has nothing to do with the person's current activity and purpose.</u>	False
Organize the kitchen and have a cup of coffee	Place the coffee cup in the dishwasher	The robot removes the cup .... If the person continue ... would likely need to retrieve the cup ..., <u>disrupting their current flow ....</u>	False

Figure 5: **Examples from the Simulator.** The figure presents examples of its future deduction and helpfulness judgment.




Figure 6: **Comparisons against human performance.** We show the Valid Task Ratio of 10 samples and the average score. The upper row is evaluated by simulation and the lower row is by labels.

**Ablation on the Search Strategy.** The results in Fig. 4 (c) compare different search strategies, where ‘CMAST-beam= $k$ ’ denotes beam search with a beam size of  $k$ . While the greedy search ( $k=1$ ) is the most efficient, it discovers only 1.4 helpful tasks on average. More time-intensive search strategies yield progressively better performance, **significantly** increasing the number of discovered tasks while keeping the accuracy ratio basically unchanged. This confirms that by scaling up the test-time thinking, our model is able to explore a broader range of future situations and cover more valuable tasks.

**Choice of Component Agents.** To investigate the influence of different component agents, we replace the agents in our framework with different LMMs, generating the following variants: CMAST-LLaVA, which is our main version, CMAST-InternVL2, which uses InternVL2-8B (Team 2024), CMAST-Qwen2, which uses Qwen2-VL-7B (Wang et al. 2024). Additionally, we report the corresponding vanilla LMMs for each variant. The results are shown in Fig. 4 (b). Comparing the same LMMs used in isolation, our framework improve the Valid Task Ratio by at least 39%. The results confirm that using different component agents consistently enhances our model’s performance, highlighting its ability to seamlessly integrate with various LMMs.

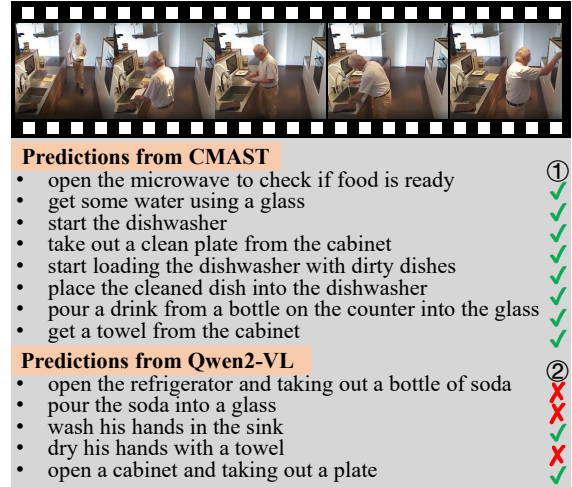


Figure 7: **The visualization** of predictions from CMAST and baseline (Wang et al. 2024). The ✓ / ✗ column indicates correct/incorrect predictions. More examples can be found in Appendix.

**Comparing with Human Performance.** We investigate whether CMAST can achieve human-level performance by randomly selecting 10 examples and asking human participants to discover tasks. The results are in Fig. 6, where the CMAST framework achieves performance comparable to the human level. We present this experiment not to claim that our method has surpassed human capabilities but to showcase its potential and the interesting phenomena observed.

## Conclusion

We introduce and formalize the problem of Human-centric Open-future Task Discovery, enabling LMMs to identify tasks that assist humans. To study it, we present HODT-Bench, a benchmark with over 2K real-world videos of diverse activities, coupled with a simulation-based protocol that evaluates open-set futures beyond observed trajectories. For stable evaluation, we additionally provide open-vocabulary task labels annotated through a semi-automated pipeline. We further propose the Collaborative Multi-Agent Search Tree framework, which leverages a multi-agent system and a scalable search tree module to structure the complex reasoning. Experiments show substantial gains in Valid Task Count and Valid Task Ratio, with consistent improvements when integrated with existing LMMs.

## Acknowledgments

This work was supported in part by the Fundamental Research Funds for Higher Education Institutions allocated to Sun Yat-sen University (Grants 25hytd007 and 2025RGZN009), in part by the Guangdong Provincial High-Level Young Talent Program (Grant 2025HYSPT0707), in part by the Tuoyuna Grant (HT-99982025-0564), in part by the Faculty Start-up Research Fund (Grant 67000-12255002), and in part by the Huawei Strategic Research Institute Talent Fund.

## References

- Aher, G. V.; Arriaga, R. I.; and Kalai, A. T. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, 337–371. PMLR.
- Ahn, M.; Dwibedi, D.; Finn, C.; Arenas, M. G.; Gopalakrishnan, K.; Hausman, K.; Ichter, B.; Irpan, A.; Joshi, N.; Julian, R.; et al. 2024. Autort: Embodied foundation models for large scale orchestration of robotic agents. *arXiv preprint arXiv:2401.12963*.
- Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Bharadhwaj, H.; Dwibedi, D.; Gupta, A.; Tulsiani, S.; Dersch, C.; Xiao, T.; Shah, D.; Xia, F.; Sadigh, D.; and Kirmani, S. 2024. Gen2Act: Human Video Generation in Novel Scenarios enables Generalizable Robot Manipulation. *arXiv preprint arXiv:2409.16283*.
- Caba Heilbron, F.; Escorcia, V.; Ghanem, B.; and Carlos Niebles, J. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 961–970.
- Cao, D.; Jia, F.; Arik, S. O.; Pfister, T.; Zheng, Y.; Ye, W.; and Liu, Y. 2024. TEMPO: Prompt-based Generative Pre-trained Transformer for Time Series Forecasting. In *International Conference on Learning Representations*.
- Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; et al. 2024. Expanding Performance Boundaries of Open-Source Multimodal Models with Model, Data, and Test-Time Scaling. *arXiv preprint arXiv:2412.05271*.
- Dai, R.; Das, S.; Sharma, S.; Minciullo, L.; Garattoni, L.; Bremond, F.; and Francesca, G. 2022. Toyota smarhome untrimmed: Real-world untrimmed videos for activity detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2533–2550.
- Damen, D.; Doughty, H.; Farinella, G. M.; Furnari, A.; Kazakos, E.; Ma, J.; Moltisanti, D.; Munro, J.; Perrett, T.; Price, W.; et al. 2022. Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision*, 1–23.
- Das, S.; Dai, R.; Koperski, M.; Minciullo, L.; Garattoni, L.; Bremond, F.; and Francesca, G. 2019. Toyota smarhome: Real-world activities of daily living. In *Proceedings of the IEEE/CVF international conference on computer vision*, 833–842.
- Driess, D.; Xia, F.; Sajjadi, M. S.; Lynch, C.; Chowdhery, A.; Ichter, B.; Wahid, A.; Tompson, J.; Vuong, Q.; Yu, T.; et al. 2023. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*.
- Grauman, K.; Westbury, A.; Byrne, E.; Chavis, Z.; Furnari, A.; Girdhar, R.; Hamburger, J.; Jiang, H.; Liu, M.; Liu, X.; et al. 2022. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18995–19012.
- Gruver, N.; Finzi, M.; Qiu, S.; and Wilson, A. G. 2023. Large language models are zero-shot time series forecasters. *Advances in Neural Information Processing Systems*, 36: 19622–19635.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hong, S.; Zheng, X.; Chen, J.; Cheng, Y.; Wang, J.; Zhang, C.; Wang, Z.; Yau, S. K. S.; Lin, Z.; Zhou, L.; et al. 2023. Metagtpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*.
- Huang, C.; Mees, O.; Zeng, A.; and Burgard, W. 2023. Visual language maps for robot navigation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 10608–10615. IEEE.
- Jia, B.; Chen, Y.; Huang, S.; Zhu, Y.; and Zhu, S.-c. 2020. LEMMA: A Multi-view Dataset for Learning Multi-agent Multi-task Activities. In *European Conference on Computer Vision*, 767–786. Springer.
- Jin, M.; Wang, S.; Ma, L.; Chu, Z.; Zhang, J.; Shi, X.; Chen, P.-Y.; Liang, Y.; Li, Y.-f.; Pan, S.; et al. 2024. TimeLLM: Time Series Forecasting by Reprogramming Large Language Models. In *International Conference on Learning Representations*.
- Katara, P.; Xian, Z.; and Fragkiadaki, K. 2024. Gen2sim: Scaling up robot learning in simulation with generative models. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 6672–6679. IEEE.
- Khandelwal, A.; Weihs, L.; Mottaghi, R.; and Kembhavi, A. 2022. Simple but effective: Clip embeddings for embodied ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14829–14838.
- Lei, J.; Yu, L.; Bansal, M.; and Berg, T. L. 2018. Tvqa: Localized, compositional video question answering. *arXiv preprint arXiv:1809.01696*.
- Liang, J.; Huang, W.; Xia, F.; Xu, P.; Hausman, K.; Ichter, B.; Florence, P.; and Zeng, A. 2023. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 9493–9500. IEEE.
- Lin, B.; Ye, Y.; Zhu, B.; Cui, J.; Ning, M.; Jin, P.; and Yuan, L. 2023a. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*.

- Lin, K.; Agia, C.; Migimatsu, T.; Pavone, M.; and Bohg, J. 2023b. Text2motion: From natural language instructions to feasible plans. *Autonomous Robots*, 47(8): 1345–1365.
- Moeslund, T. B.; Hilton, A.; and Krüger, V. 2006. A survey of advances in vision-based human motion capture and analysis. *Computer vision and image understanding*, 104(2-3): 90–126.
- OpenAI. 2024. Learning to Reason with LLMs. <https://openai.com/index/learning-to-reason-with-llms/>.
- OpenAI. 2025. OpenAI o3-mini System Card. Technical report, OpenAI. Accessed: 2025-03-04.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36: 53728–53741.
- Sigurdsson, G. A.; Varol, G.; Wang, X.; Farhadi, A.; Laptev, I.; and Gupta, A. 2016. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, 510–526. Springer.
- Soomro, K. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Team, O. 2024. InternVL2: Better than the Best—Expanding Performance Boundaries of Open-Source Multimodal Models with the Progressive Scaling Strategy. <https://internvl.github.io/blog/2024-07-02-InternVL-2.0/>.
- Team, Q. 2025. Qwen2.5-VL.
- Wang, L.; Ling, Y.; Yuan, Z.; Shridhar, M.; Bao, C.; Qin, Y.; Wang, B.; Xu, H.; and Wang, X. 2023a. Gensim: Generating robotic simulation tasks via large language models. *arXiv preprint arXiv:2310.01361*.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; et al. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Wang, Y.; Xian, Z.; Chen, F.; Wang, T.-H.; Wang, Y.; Fragkiadaki, K.; Erickson, Z.; Held, D.; and Gan, C. 2023b. Robogen: Towards unleashing infinite data for automated robot learning via generative simulation. *arXiv preprint arXiv:2311.01455*.
- Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Zhang, S.; Zhu, E.; Li, B.; Jiang, L.; Zhang, X.; and Wang, C. 2023. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Yang, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.; Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Liu, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; Guo, Z.; and Fan, Z. 2024a. Qwen2 Technical Report. *arXiv:2407.10671*.
- Yang, S.; Luo, Q.; Pani, A.; and Yang, Y. 2024b. BBSEA: An Exploration of Brain-Body Synchronization for Embodied Agents. *arXiv preprint arXiv:2402.08212*.
- Yang, Y.; Zhou, T.; Li, K.; Tao, D.; Li, L.; Shen, L.; He, X.; Jiang, J.; and Shi, Y. 2024c. Embodied multi-modal agent trained by an llm from a parallel textworld. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26275–26285.
- Yang, Z.; Liu, J.; Chen, P.; Cherian, A.; Marks, T. K.; Le Roux, J.; and Gan, C. 2024d. RILA: Reflective and Imaginative Language Agent for Zero-Shot Semantic Audio-Visual Navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16251–16261.
- Yu, W.; Gileadi, N.; Fu, C.; Kirmani, S.; Lee, K.-H.; Arenas, M. G.; Chiang, H.-T. L.; Erez, T.; Hasenclever, L.; Humplik, J.; et al. 2023. Language to rewards for robotic skill synthesis. *arXiv preprint arXiv:2306.08647*.
- Yu, Y.; Yao, Z.; Li, H.; Deng, Z.; Jiang, Y.; Cao, Y.; Chen, Z.; Suchow, J.; Cui, Z.; Liu, R.; et al. 2025. Fincon: A synthesized llm multi-agent system with conceptual verbal reinforcement for enhanced financial decision making. *Advances in Neural Information Processing Systems*, 37: 137010–137045.
- Zellers, R.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6720–6731.
- Zeng, R.; Chen, X.; Liang, J.; Wu, H.; Cao, G.; and Guo, Y. 2024. Benchmarking the Robustness of Temporal Action Detection Models Against Temporal Corruptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18263–18274.
- Zhang, Y.; Li, B.; Liu, h.; Lee, Y. j.; Gui, L.; Fu, D.; Feng, J.; Liu, Z.; and Li, C. 2024. LLaVA-NeXT: A Strong Zero-shot Video Understanding Model.
- Zhao, X.; Weber, C.; and Wermter, S. 2024. Agentic Skill Discovery. *arXiv preprint arXiv:2405.15019*.
- Zhou, H.; Martín-Martín, R.; Kapadia, M.; Savarese, S.; and Niebles, J. C. 2023. Procedure-aware pretraining for instructional video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10727–10738.
- Zhou, Z.; Atreya, P.; Lee, A.; Walke, H.; Mees, O.; and Levine, S. 2024. Autonomous improvement of instruction following skills via foundation models. *arXiv preprint arXiv:2407.20635*.